

Collecting Image Description Datasets using Crowdsourcing

Ramakrishna Vedantam, C. Lawrence Zitnick and Devi Parikh

Virginia Tech
Microsoft Research

Abstract. We describe our two new datasets with images described by humans. Both the datasets were collected using Amazon Mechanical Turk, a crowdsourcing platform. The two datasets contain significantly more descriptions per image than other existing datasets. One is based on a popular image description dataset called the UIUC Pascal Sentence Dataset, whereas the other is based on the Abstract Scenes dataset containing images made from clipart objects. In this paper we describe our interfaces, analyze some properties of and show example descriptions from our two datasets.

Keywords: Image Description, Vision

1 Introduction

Recent works have explored the connection between Natural Language and Images. There is particular interest in both the Vision and NLP communities to explore the common ground between the two areas. On the vision side, understanding the interplay with language can help drive vision systems that communicate with humans, summarize important aspects in the scene etc. On the language side, there is much interest in grounding language learning with perceptual cues.

To facilitate and spur future progress in these areas, appropriate datasets are critical. A lot of datasets of image descriptions exist [1–5]. However, the most number of sentences collected by any dataset so far is five. We introduce two image-description datasets with 50 captions for every image. We call these datasets ABSTRACT-50S and PASCAL-50S. Our two datasets are “gold-standard” in the sense that the sentences are all written by human subjects with the intention of describing the image. The ABSTRACT-50S dataset is based on the dataset of Zitnick and Parikh [1] which has cartoon-like abstract images. This dataset, synthetically generated using crowdsourcing, provides opportunities to focus on image-semantics without the inhibition of (still) noisy visual detectors. The second dataset, is based on images from the UIUC Pascal Sentence Dataset. These are real images collected from Flickr.

The UIUC Pascal Sentence Dataset has been used in various works for describing images [6, 7], doing better semantic segmentation [8], understanding properties of image descriptions [9], etc. Other works have leveraged Abstract Images for performing zero-shot learning [10], generating images from text [11],



Please describe what is going on in this image.

Sentence:

Fig.1: A snapshot of our sentence collection interface, shown to subjects on Amazon Mechanical Turk.

and understanding the semantics of images [1]. We hope our new datasets will facilitate further progress along these varied directions.

2 Related Work

The most popular Image-Sentence dataset is the UIUC Pascal Sentence Dataset [2]. This dataset contains 5 human written descriptions for 1000 images. The SBU captioned photo dataset [3] contains one description per image for a million images, mined from the web. This dataset has automatically mined descriptions, which are not “gold-standard”. That is, there exist many descriptions which are not relevant to the image content [4]. Recent works have looked at the problem of identifying “visual” text [12]. Further progress could lead to image description datasets with a large number of images. Our focus, in contrast, is to create a dataset that captures fine grained notions of object importance and description styles. Thus we create a dataset with a large number of descriptions per Image. The Abstract Scenes dataset contains cartoon like images with two descriptions. The recently released MS-COCO dataset [13] contains five sentences for a collection of over 100k images. The Flickr8k dataset contains five descriptions for a collection of 8000 images. The images in this dataset were queried for actions. Subjects were instructed to describe the major actions and objects in the scene. Other datasets of images and associated descriptions includes the Image Clef Dataset, which tends to have longer sentences (21 words) [5]. We describe two new datasets. First is the PASCAL-50S dataset where we collect 50 sentences



(a) A sample image shown to workers

This is a whacky image.

Reasons for rejection:

- (1) Does not actually describe what is going on in the scene.
- (2) Too short.

A man is next to a woman, the man is on the couch, the woman is on the couch, a dog is on the man

Reasons for rejection:

- (1) Too literal.
- (2) Others probably would not describe this scene the same way.

(b) Examples of sentences that would be rejected, with reasons

Fig. 2: Explanation of our rejection criteria. Our goal was to collect sentences representative of image content

per image for the 1000 images from UIUC Pascal Sentence dataset. The second is the ABSTRACT-50S dataset where we collect 50 sentences for a subset of 500 images from the Abstract Scenes dataset.

3 Interface

Our goal was to collect image descriptions that are objective and representative of the image content. We first showed subjects a set of images on mechanical turk and asked them to “describe” them. We found that when asked to describe images, subjects would use their imagination and often not produce descriptions that are relevant to image content. We thus asked the subjects to “transcribe” the major aspects of the scene into descriptions. We found that making this change helped elicit more objective and image-related descriptions from the Subjects. The exact details of our interface are as follows. Subjects are shown an image and a text box, asking them to describe what is “going on” in the image. We instruct the workers that good transcriptions are those that others are also likely to provide (see Fig. 2). This includes writing descriptions rather than “dialogs” or overly descriptive sentences. They were encouraged to capture the main aspects of the scene. Subjects were told that a good description should help others recognize the image from a collection of similar images. Instructions also mentioned that work with poor grammar would be rejected. Snapshots of our interface can be seen in Fig. 1. Overall, we had 465 subjects for ABSTRACT-50S and 683 subjects for PASCAL-50S datasets. We ensure that each sentence for an image is written by a different subject.

To ensure that the sentences are of desired quality, certain qualification criteria were imposed. Subjects were required to be from the United States. Only subjects who had a 95% HIT (Human Intelligence Task) approval rate and had been approved 500 times were considered eligible on Amazon Mechanical Turk.

4 Analysis and Results

Overall, we find that the sentences collected for PASCAL-50S are on average 8.8 words in length. In contrast, on the ABSTRACT-50S dataset we find that the description length is 10.59 words. This could be because of the tighter semantic sampling that the abstract images impose. Thus, sentences tend to be more detailed to be discriminative. We show some scenes from our PASCAL-50S dataset in Fig. 3 and some from our ABSTRACT-50S dataset in Fig. 4 respectively.



The scene of a poor neighbor hood
 A run down house needs some repairs in town
 A run down building as flora growing on and in it
 An older stucco house with plants growing on it
 Small, dilapidated house that is falling apart
 Ivy grows from abandoned motel
 A small house with a detached door



The shot of several pointy things in the middle of the city
 Several pointy structures line the horizon
 several unique structures are shown in the background
 A car is passing lookout towers
 Three spires rise up from behind the highway
 These buildings were spaceships in men in black
 A car is driving by the Kuwait towers
 Three pointy towers with some designs

Fig. 3: Sample images with a subset of collected sentences from PASCAL-50S. Notice the rich variation in descriptions, because of collecting a large number of descriptions per image

5 Conclusions

In this paper, we describe two new datasets ABSTRACT-50S and PASCAL-50S with 50 sentences per image. We provide interface details and a background on the motivation for these datasets. These proposed datasets capture the many ways in which humans describe images. We hope these two new datasets will spur further research on exploring the connection between vision and language, two primary interaction modalities for humans, and lead to future research in building more intelligent systems.



Jenny is sad that she dropped her hamburger.

Jenny is sad because she dropped her hamburger on the ground.

Jenny is upset for dropping her hamburger on the ground at a BBQ.

Jenny dropped her hamburger on the ground because her sunglasses were too dark and she couldn't see the table.

Jenny looks unhappy as she sees the hamburger on the ground.

Jenny is unhappy that the hamburger fell off the grill.



Mike and Jenny run away from the snake.

Jenny and Mike run away from a snake while an owl watches from a tree.

A snake chases away Mike and Jenny from the swing set.

Mike and Jenny are running away from a snake on a sunny day.

The snake scared Mike and Jenny and they ran away.

Jenny and Mike are running and Mike is being chased by a snake.

Fig. 4: Sample images with a subset of collected sentences from ABSTRACT-50S. Notice the rich variation in descriptions, because of collecting a large number of descriptions per image

References

1. Zitnick, C.L., Parikh, D.: Bringing semantics into focus using visual abstraction. In: CVPR. (2013) [1](#), [2](#)
2. Rashtchian, C., Young, P., Hodosh, M., Hockenmaier, J.: Collecting image annotations using amazon’s mechanical turk. In: Proceedings of the NAACL HLT 2010 Workshop on Creating Speech and Language Data with Amazon’s Mechanical Turk. CSLDAMT ’10, Stroudsburg, PA, USA, Association for Computational Linguistics (2010) 139–147 [1](#), [2](#)
3. Ordonez, V., Kulkarni, G., Berg, T.L.: Im2text: Describing images using 1 million captioned photographs. In: Neural Information Processing Systems (NIPS). (2011) [1](#), [2](#)
4. Hodosh, M., Young, P., Hockenmaier, J.: Framing image description as a ranking task: Data, models and evaluation metrics. J. Artif. Intell. Res. (JAIR) **47** (2013) 853–899 [1](#), [2](#)
5. Mller, H., Clough, P., Deselaers, T., Caputo, B.: ImageCLEF: Experimental Evaluation in Visual Information Retrieval. 1st edn. Springer Publishing Company, Incorporated (2010) [1](#), [2](#)
6. Kulkarni, G., Premraj, V., Dhar, S., Li, S., Choi, Y., Berg, A.C., Berg, T.L.: Baby talk: Understanding and generating simple image descriptions. In: Proceedings of the 2011 IEEE Conference on Computer Vision and Pattern Recognition. CVPR ’11, Washington, DC, USA, IEEE Computer Society (2011) 1601–1608 [1](#)
7. Mitchell, M., Han, X., Hayes, J.: Midge: Generating descriptions of images. In: Proceedings of the Seventh International Natural Language Generation Conference. INLG ’12, Stroudsburg, PA, USA, Association for Computational Linguistics (2012) 131–133 [1](#)
8. Fidler, S., Sharma, A., Urtasun, R.: A sentence is worth a thousand pixels. In: CVPR, IEEE (2013) 1995–2002 [1](#)
9. Berg, A.C., Berg, T.L., III, H.D., Dodge, J., Goyal, A., Han, X., Mensch, A., Mitchell, M., Sood, A., Stratos, K., Yamaguchi, K.: Understanding and predicting importance in images. In: CVPR, IEEE (2012) 3562–3569 [1](#)
10. Antol, S., Zitnick, C.L., Parikh, D.: Zero-Shot Learning via Visual Abstraction. In: ECCV. (2014) [1](#)
11. Zitnick, C.L., Parikh, D., Vanderwende, L.: Learning the visual interpretation of sentences. In: The IEEE International Conference on Computer Vision (ICCV). (December 2013) [1](#)
12. Kuznetsova, P., Ordonez, V., Berg, A.C., Berg, T.L., Choi, Y.: Generalizing image captions for image-text parallel corpus. In: ACL (2), The Association for Computer Linguistics (2013) 790–796 [2](#)
13. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft COCO: Common objects in context. In: European Conference on Computer Vision (ECCV). (2014) [2](#)